

Decision Making in Intelligent Agents

Mats Danielson

Stockholm University, Sweden & Royal Institute of Technology, Sweden

Love Ekenberg

Stockholm University, Sweden & Royal Institute of Technology, Sweden

INTRODUCTION


There are several ways of building complex distributed software systems, for example in the form of software agents. But regardless of the form, there are some common problems having to do with specification contra execution. One of the problems is the inherent dynamics in the environment many systems are exposed to. The properties of the environment are not known with any precision at the time of construction. This renders a specification of the system incomplete by definition. A traditional **software agent** is only prepared to handle situations conceived of and implemented at compile-time. Even though it can operate in varying contexts, its decision making abilities are static. One remedy is to prepare the distributed components for a truly dynamic environment, i.e. an environment with changing and somewhat unpredictable conditions. A rational **software agent** needs both a representation of a decision problem at hand and means for evaluation. AI has traditionally addressed some parts of this problem such as representation and reasoning, but has hitherto to a lesser degree addressed the decision making abilities of independent distributed software components (Ekenberg, 2000a, 2000b). Such decision making often has to be carried out under severe uncertainty regarding several parameters. Thus, methods for independent **decision making** components should be able to handle uncertainties on the probabilities and utilities involved. They have mostly been studied as means of representation, but are now being developed into functional theories of decision making suitable for dynamic use by software agents and other dynamic distributed components. Such a functional theory will also benefit analytical decision support systems intended to aid humans in their decision making. Thus, the generic term *agent* below stands for a dynamic software component as well as a human or a group of humans assisted by intelligent software.

BACKGROUND

Ramsey (1926/78) was the first to suggest a theory that integrated ideas on subjective probability and utility in presenting (informally) a general set of axioms for preference comparisons between acts with uncertain outcomes (probabilistic decisions). von Neumann and Morgenstern (1947) established the foundations for a modern theory of utility. They stated a set of axioms that they deemed reasonable to a rational decision-maker (such as an agent), and demonstrated that the agent should prefer the alternative with the highest expected utility, given that she acted in accordance with the axioms. This is the principle of maximizing the **expected utility**. Savage (1954/72) published a thorough treatment of a complete theory of subjective expected utility. Savage, von Neumann, and others structured decision analysis by proposing reasonable principles governing decisions and by constructing a theory out of them. In other words, they (and later many others) formulated a set of axioms meant to justify their particular attitude towards the utility principle, cf., e.g., Herstein and Milnor (1953), Suppes (1956), Jeffrey (1965/83), and Luce and Krantz (1971). In classical **decision analysis**, of the types suggested by Savage and others, a widespread opinion is that utility theory captures the concept of rationality.

After Raiffa (1968), probabilistic decision models are nowadays often given a tree representation (see Fig. 1). A **decision tree** consists of a root, representing a decision, a set of event nodes, representing some kind of uncertainty and consequence nodes, representing possible final outcomes. In the figure, the decision is a square, the events are circles, and final consequences are triangles. Events unfold from left to right, until final consequences are reached. There may also be more than one decision to make, in which case the sub-decisions are made before the main decision.

Figure 1. Decision tree



In decision trees, probability distributions are assigned in the form of weights (numbers) in the probability nodes as measures of the uncertainties involved. Obviously, such a numerically precise approach puts heavy demands on the input capability of the agent. The shortcomings of this representation are many, and have to be compensated for, see, e.g., (Ekenberg, 2000a). Among other things, the question has been raised whether people are capable of providing the input information that utility theory requires (cf., e.g., (Fischhoff et al., 1983)). For instance, most people cannot clearly distinguish between probabilities ranging roughly from 0.3 to 0.7 (Shapira, 1995). Similar problems arise in the case of artificial agents, since utility-based artificial agents usually base their reasoning on human assessments, for instance in the form of induced preference functions. The so-called reactive agents, for which this does not hold true, have not been put to use in dynamic domains involving uncertainty (cf., e.g., (Russell & Norvig, 1995)). Furthermore, even if an agent would be able to discriminate between different probabilities, very often complete, adequate, and precise information is missing.

Consequently, during recent years of rather intense research activities several alternative approaches have emerged. In particular, first-order approaches, i.e., based on sets of probability measures, upper and lower probabilities, and **interval probabilities**, have prevailed. A main class of such models has been focused on expressing probabilities in terms of intervals. In 1953, the concept of capacities was introduced (Choquet, 1953/54). This representation approach was further

developed in (Huber, 1973, Huber & Strassen, 1973). Capacities have subsequently been used for modelling imprecise probabilities as intervals (capacities of order 2 (Denneberg, 1994)). Since the beginning of the 1960s the use of **first-order** (interval-valued) probability functions, by means of classes of probability measures, has been integrated in classical probability theory by, e.g., Smith (1961) and Good (1962). Similarly, Dempster (1967) investigated a framework for modelling upper and lower probabilities, which was further developed by Shafer (1976), where a representation of belief in states or events was provided. Within the AI community the Dempster-Shafer approach has received a good deal of attention. However, their formalism seems to be too strong to be an adequate representation of belief (Weichselberger & Pöhlman, 1990).

Other representations in terms of upper and lower probabilities have been proposed by, i.a., Hodges and Lehmann (1952), Hurwicz (1951), Wald (1950), Kyburg (1961), Levi (1974, 1980), Walley (1991), Danielson and Ekenberg (1998, 2007), and Ekenberg et al. (2001). Upper and lower previsions have also been investigated by various authors. For instance, Shafer et al. (2003) suggests a theory for how to understand subjective probability estimates based on Walley (1991). A few approaches have also been based on logic, e.g., Nilsson (1986). He develops methods for dealing with sentences involving upper and lower probabilities. This kind of approaches has been pursued further by, among others, Wilson (1999).

4 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage: www.igi-global.com/chapter/decision-making-intelligent-agents/10283

Related Content

Finding Time for Creativity in Higher Education Writing Through Generative AI

K. M. Begian-Lewis and Colleen J. Hart (2024). *Impacts of Generative AI on Creativity in Higher Education* (pp. 47-68).

www.irma-international.org/chapter/finding-time-for-creativity-in-higher-education-writing-through-generative-ai/355427

Automatic Topic Ontology Construction Using Semantic Relations from WordNet and Wikipedia

V. Subramaniaswamy (2013). *International Journal of Intelligent Information Technologies* (pp. 61-89).

www.irma-international.org/article/automatic-topic-ontology-construction-using-semantic-relations-from-wordnet-and-wikipedia/93153

Exploring Multi-Path Communication in Hybrid Mobile Ad Hoc Networks

Roberto Speicys Cardoso and Mauro Caporuscio (2010). *International Journal of Ambient Computing and Intelligence* (pp. 1-12).

www.irma-international.org/article/exploring-multi-path-communication-hybrid/47173

AI-Driven Threat Detection and Prevention: Enhancing Cybersecurity with Machine Learning and Predictive Intelligence

Tarun Kumar Vashishth, Vikas Sharma, Sangeeta Sharma, Mukesh Kumar Sharma, Rajeev Sharma, Sachin Chaudhary and Aarushi Vashishth (2026). *Implementing Enterprise Cybersecurity With AI* (pp. 57-84).

www.irma-international.org/chapter/ai-driven-threat-detection-and-prevention/395160

Construction of Domain Ontologies: Sourcing the World Wide Web

Jongwoo Kim and Veda C. Storey (2011). *International Journal of Intelligent Information Technologies* (pp. 1-24).

www.irma-international.org/article/construction-domain-ontologies/54064